arXiv:1409.2081v1 [cs.GR] 7 Sep 2014

History-free Collision Response for Deformable Surfaces

Juntao Ye*
Institute of Automation, Chinese Academy of Sciences

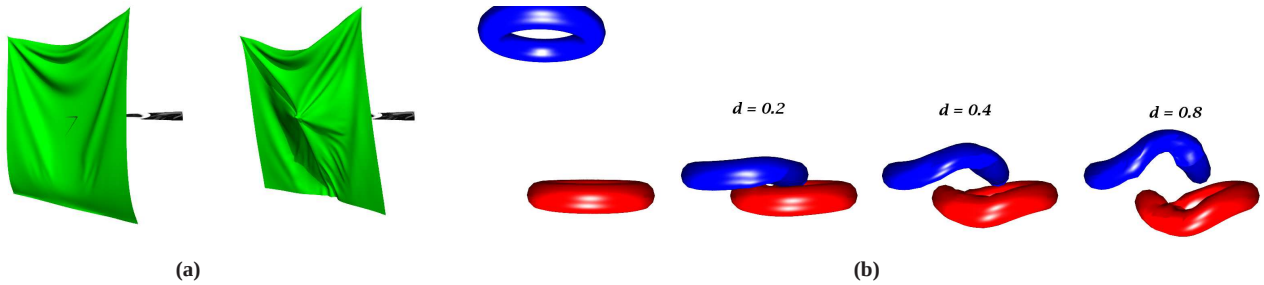

(a) (b)

Figure 1: *(a) An existing penetration configuration is recovered with our DCD-based response method. (b) The blue torus is falling on to the red one. By setting different post-response distances, a range of behaviors from inelastic to elastic collision can be simulated.*

Abstract

Continuous collision detection (CCD) and response methods are widely adopted in dynamics simulation of deformable models. They are history-based, as their success is strictly based on an assumption of a collision-free state at the start of each time interval. On the other hand, in many applications surfaces have normals defined to designate their orientation (i.e. front- and back-face), yet CCD methods are totally blind to such orientation identification (thus are orientation-free). We notice that if such information is utilized, many penetrations can be untangled. In this paper we present a history-free method for separation of two penetrating meshes, where at least one of them has clarified surface orientation. This method first computes all edge-face (E-F) intersections with discrete collision detection (DCD), and then builds a number of penetration stencils. On response, the stencil vertices are relocated into a penetration-free state, via a global displacement minimizer. Our method is very effective for handling penetration between two meshes, being it an initial configuration or in the middle of physics simulation. The major limitation is that it is not applicable to self-collision within one mesh at the time being.

CR Categories: I.3.3 [Computer Graphics]: Three-Dimensional Graphics and Realism—Display Algorithms I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Radiosity;

Keywords: collision response, deformable models, cloth simulation

1 Introduction

Collisions are often unavoidable in modeling deformable surfaces such as cloth. The majority simulation systems adopt CCD to predict impending collisions, then attempt to prevent them from happening by altering particles' velocities. The success of CCD-based response method relies on a hard constraint: a collision-free state for not only the initial configuration but also the start of every time interval in the simulation. This requirement makes CCD methods to be called history-based. A major advantage of CCD methods is that they do not rely on surface orientation information to decide whether a geometric primitive is located on the incorrect side of a surface or not. This is very important for surfaces (e.g. a sheet model) which typically do not present orientation information. However, there also exists surfaces which do exhibits orientation information, either explicitly or implicitly. For example, in a polygonal sphere model, the polygon normals usually point to the outside of the sphere; another example is the simulation of a dressed avatar, in which the normals of the body surface also points to the outside so that no cloth vertices are allowed to go inside the body. We call them *oriented surfaces*. CCD methods, designed to handle generic collision context which could involve either oriented or un-oriented surfaces, are totally blind to such useful orientation information. Nevertheless, normals are indispensable ingredients in CCD and their directions for colliding polygons are so picked as to be consistent with the negative approaching vertex velocities. In this sense orientations of those polygons are derived dynamically from the history. For oriented surfaces, this derivation seems redundant. An interesting question is, if orientation is available, will it be more helpful than history information?

Our goal is to put forward a new method that takes advantage of surface orientation if it is available, and no longer relies on the history information. Different from CCD as prevention-based methods, our method embraces the *repair* strategy – penetrations are allowed to occur but will be detected and penetrating objects will be separated. It employs *discrete collision detection* (DCD) so it is history-free. DCD only finds out edge-face (E-F) intersections for two given meshes, thus is faster than CCD for two points: bounding volumes are more tight-fitting and no cubic solver is needed. Please note that at the time being our method is capable of handling collisions between two different meshes, at least one of them being oriented. How to extend our method to handle self-collision is still under investigation.

*e-mail:juntao.ye@ia.ac.cn

2 Related work

A number of research groups have worked on the CCD method. [Provot 1997] proposed a cubic solver for detecting penetration of two elastic moving triangles. [Bridson et al. 2002] presents a comprehensive way of preventing intersections to occur. CCD once suffered a serious problem that the cubic polynomial solver is vulnerable to round-off error. This problem was reported to be solved not long ago [Brochu et al. 2012]. CCD helped the asynchronous contact model [Harmon et al. 2009; Ainsley et al. 2012] to produce very realistic effects. In CCD, every detected collision event carries a time tag, but they are often ignored, as resolving all the events strictly according to their time order can be too slow to halt the simulation. Thus all events are often assumed to occur at exactly the same time instant and simultaneous response is adopted [Harmon et al. 2008]. This demands the time interval for CCD to be of moderate size.

CCD will suffer even serious problem in applications in which an intersection-free initial state is impossible, or the simulation context is subject to external constraints forcing the deformable surface into illegal states for a number of consecutive steps. In these situations if some vertices end up on the back-face side of another surface, they will be "remembered" as not, and the simulator will work hard to keep them on the wrong side forever. Therefore, a "repair" mechanism is needed for robust simulation. Baraff *et. al* [Baraff et al. 2003] were probably the first to address the necessity of collision correction in complex simulation environment, and put forward a partial solution. Their method is limited to closed regions and cannot handle intersection involving mesh boundary. Wicke *et. al* [Wicke et al. 2006] further investigated this problem by taking consideration of boundary collisions. They analyzed all possible collision configurations and presented a classification criterion for intersection contours. An intersection contour minimization (ICM) method [Volino and Magnenat-Thalmann 2006] was proposed. While their algorithm are capable of handling certain very complex situations, it also fails for some other situations. The failure or slow convergence of ICM is due to ambiguous normal direction used according to [Ye and Zhao 2012], therefore it is only capable of handling boundary-involved cases but is not good at closed contours. Narain *et. al* [Narain et al. 2012] encountered the interpenetration problem while re-meshing cloth to enrich details, and the penetrating vertices are moved back in a manner quite similar to our method in this paper.

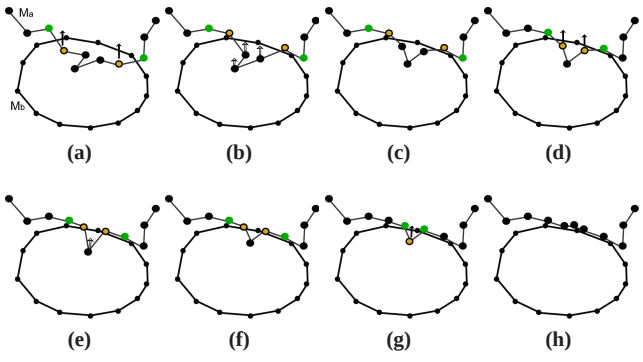

Figure 2: *The 2D illustration of how our method works. Mesh M_b is oriented surface of a solid object, and M_a is un-oriented deformable surface. In the penetration configuration: vertex of an intersecting edge is marked green if it is in legal state, or marked yellow if it is in illegal state; a vertex is marked black if it is not on any intersecting edge.*

3 Algorithm overview

We show how our method works by a 2D illustration of a penetration configuration between one oriented surface M_b and one un-oriented surface M_a in Fig. 2. The illegally located vertices are iteratively detected and then relocated. First the DCD algorithm finds all E-F intersections, in which the face is from the oriented face and the edge is from the un-oriented one. Then the legality of each end-vertex of these edges is then determined according to the front-face identification: legal vertices are marked green and illegal vertices are yellow. For those edges that do not intersect any faces, the legality of their end-vertices can not be determined and these vertices are marked black (Fig. 2a). The illegal (yellow) vertices are to be relocated to legal positions via our *minimum positional displacement* algorithm (Fig. 2b), and also the displacement vector will be diffused to the vicinity (Fig. 2c). These yellow vertices become green in the next round of DCD (Fig. 2d). Repeat this process until there is no E-F intersections (Fig. 2e -2h). Our algorithm is so designed that the yellow vertices are moved in parallel, along the shortest paths, to reach penetration-free state.

4 Minimum positional displacement

In this paper, the surfaces are described as polygonal (particularly triangular) meshes. The entire mesh, at any time instant, can be viewed as a point in a high-dimensional space. For an *n*-vertex mesh, we denote this point as $\mathbf{x}(t) = [\mathbf{x}_0(t), \mathbf{x}_1(t), \cdots, \mathbf{x}_{n-1}(t)] \in \mathbb{R}^{3n}$. At any time instant, surface intersections consist of a number of face-face (F-F) intersections.

Fig. 3(a)(b) shows two cases of F-F penetration, in which one face has a normal designating its front-face orientation and the other is un-oriented. In both cases, vertex $\mathbf{x}_0$ is located illegally on the back-face side of $\mathbf{x}_1\mathbf{x}_2\mathbf{x}_3$ (Fig 3(c)). If we could relocate these vertices so that $\mathbf{x}_0$ goes to the other side of the face (Fig 3(d)), this penetration will be resolved. As we can see an F-F penetration can always be decomposed into two E-F penetrations, the E-F penetration then becomes the elementary penetration to be resolved. We borrow a terminology from [Harmon et al. 2008] and call the quadruple $\{\mathbf{x}_0, \mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3\}$ a *penetration stencil*. A stencil has three vertices coming from an oriented face, and one vertex from an edge, be it from an oriented surface or not. The F-F penetration in Fig 3a has two E-F penetrations, leading to two stencils (yet they are identical). The two E-F penetrations in Fig 3b result in only one stencil, as one face has no orientation. Please note if both faces are orientated we interchangeably regarded one as oriented and other as un-oriented in the iterations.

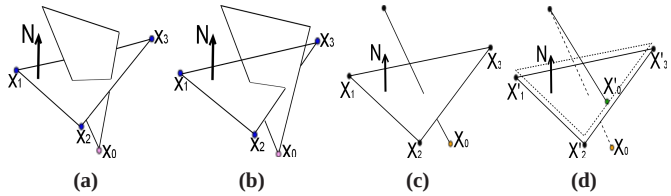

Figure 3: *(a)(b) are two types of face-face intersection; (c) four vertices $\{\mathbf{x}_0, \mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3\}$ form a V-F penetration stencil, and in (d) penetration is resolved by relocating four vertices so that $\mathbf{x}'_0$ is above the triangle plane.*

Penetration Constraints. We view each E-F penetration as the violation of a scalar-valued constraint function, $D(\mathbf{x})$, with $D(\mathbf{x}) < 0$ whenever $\mathbf{x}$ is in a penetration configuration. For V-F penetration the constraint function can be considered as a *signed distance* from

vertex $\mathbf{x}_0$ to face $\mathbf{x}_1\mathbf{x}_2\mathbf{x}_3$:

$$D(\mathbf{x}) = \mathbf{n} \cdot [\mathbf{x}_0 - (\alpha_1\mathbf{x}_1 + \alpha_2\mathbf{x}_2 + \alpha_3\mathbf{x}_3)] \quad (1)$$

where $\mathbf{n}$ is the triangle normal, and $\alpha_1, \alpha_2, \alpha_3$ are the barycentric coordinates of any point on the plane spanned by the triangle. As will be clarified later on, we choose $\alpha_{1,2,3} = 1/3$, anticipating that the three vertices are undergoing the same displacement vector during the relocation.

While each constraint depends on only four vertices, it can be differentiated with respect to $\mathbf{x}$ to yield the *constraint gradient*, ∇D, a row vector in configuration space. For a V-F penetration, the constraint gradient expressed in the local indices of the stencil is

$$\nabla D = (\mathbf{n}^T, -\frac{1}{3}\mathbf{n}^T, -\frac{1}{3}\mathbf{n}^T, -\frac{1}{3}\mathbf{n}^T). \quad (2)$$

To elevate the constraint gradient in the $\mathbb{R}^{3n}$ space, we simply map the local indices to their global position and re-write ∇D as a $1\times 3n$ row vector, with zeros everywhere else, in the style of finite element stiffness matrix assembly. Please note that the inner product $\nabla D\mathbf{x}$ is also the signed distance, thus $\nabla D\mathbf{x} = D(\mathbf{x})$.

Displacement for a single constraint. Let unprimed and primed quantities refer to the *pre-* and *post*-response states, respectively. For example, $\mathbf{x}'$ is the post-response position. A valid response satisfies two properties. First, it pushes colliding objects apart by relocating each of the four vertices along the constraint direction, *i.e.*, $\mathbf{x}' = \mathbf{x} + \mathbf{M}^{-1}\nabla D^T\lambda$ for some (unknown) non-negative multiplier λ. The diagonal mass matrix $\mathbf{M}$ here is to meet the requirement of conservation law, thus $\mathbf{M}^{-1}\nabla D^T\lambda$ is the displacement vector for all four vertices. Second, the post-response signed distance should be non-negative: $\nabla D'\mathbf{x}' \geq 0$.

For properly triangulated mesh used in dynamics simulation, assuming equal mass for all vertices is very common. In this case, together with the carefully chosen $\alpha_{1,2,3} = 1/3$, the normal $\mathbf{n}$ of face $\mathbf{x}_1\mathbf{x}_2\mathbf{x}_3$ remains unchanged after response. Thus there is $\nabla D' = \nabla D$ and it is safe to use $\nabla D\mathbf{x}'$ as the post-response signed distance. For an inelastic collision, we seek the maximal dissipation among all valid responses, per the definition of purely inelastic. Thus we minimize the post-response signed distance: $\nabla D'\mathbf{x}' = 0$, which is equivalent to

$$\nabla D\mathbf{x}' = \nabla D(\mathbf{x} + \mathbf{M}^{-1}\nabla D^T\lambda) = 0,$$

from which λ can be computed.

Displacements for multiple constraints. We are typically faced with k penetration constraints, *i.e.*, k V-F penetrations with possibly non-disjoint stencils. Now let $\nabla D = [\nabla D_1^T, \cdots, \nabla D_k^T]^T$ be a $k \times 3n$ matrix whose rows span the possible displacement directions. For any vector $\lambda \in \mathbb{R}^k$, $\mathbf{x}' = \mathbf{x} + \mathbf{M}^{-1}\nabla D^T\lambda$ corresponds to the application of a linear combination of correcting displacements to $\mathbf{x}$. As in the single-constraint case, we require that $\lambda_1 \cdots \lambda_k$ be non-negative, since displacements are to cancel the negative distance, and lead to a non-negative post-response distance, *i.e.*, that every row of $\nabla D'\mathbf{x}' = [\nabla D_1'\mathbf{x}', \cdots, \nabla D_k'\mathbf{x}']^T$ be non-negative.

We propose an algorithm for approximating λ by assuming the post-response distance to be exactly zero: $\nabla D'\mathbf{x}' = \mathbf{0}$. As evaluating $\nabla D'$ requires computation of the triangle normals according to their new (unknown) vertex positions, we instead use ∇D to approximate $\nabla D'$ and instead demand $\nabla D\mathbf{x}' = \mathbf{0}$. We relax the conditions on a response being valid to allow both positive and negative entries in λ, then $\mathbf{x}'$ is the minimizer of

$$||\mathbf{x}' - \mathbf{x}||_M^2, \quad \text{subject to} \quad \nabla D\mathbf{x}' = \mathbf{0}. \quad (3)$$

This minimization seeks new vertex positions that are as close as possible to the old positions, in the sense that position differences are penalized by the corresponding particle masses. This minimization projects the illegal vertices onto the orthogonal complement of the span of the columns of ∇D^T. We may repose the above as an extremization of the augmented functional

$$W(\mathbf{x}', \lambda) = \frac{1}{2}||\mathbf{x}' - \mathbf{x}||_M^2 + (\nabla D\mathbf{x}')^T\lambda, \quad (4)$$

with respect to $(\mathbf{x}', \lambda)$, where λ is a vector of Lagrange multiplier. The corresponding stationary equations are

$$\mathbf{0} = \frac{\partial W}{\partial \mathbf{x}'} = \mathbf{x}' - \mathbf{x} + \mathbf{M}^{-1}\nabla D^T\lambda, \quad (5)$$

$$\mathbf{0} = \frac{\partial W}{\partial \lambda} = \nabla D\mathbf{x}'. \quad (6)$$

The first stationary equation guarantees that the response acts only along the ∇D direction, and the second one ensures vanishing negative distance. Substituting the first into the second yields

$$\nabla D\mathbf{M}^{-1}\nabla D^T\lambda = \nabla D\mathbf{x} \quad (7)$$

for λ, and then substituting it into the first stationary equation to obtain positional displacement for a set of k simultaneous collisions.

Elastic collision. If we don't have to pursue purely inelastic collision, the post-response signed distance can be set to certain positive value, i.e., $\nabla D\mathbf{x}' = \mathbf{d}$, with $\mathbf{d} = (d_1, d_2, \cdots, d_k)^T$. In this case, Equ. 7 becomes

$$\nabla D\mathbf{M}^{-1}\nabla D^T\lambda = \nabla D\mathbf{x} - \mathbf{d}. \quad (8)$$

For models that assume isotropic physical properties across the whole mesh, we can simply set identical values for d_1, d_2, $\cdots$, d_k.

Implementation. After performing DCD on two meshes, we may obtain a number of penetration stencils. To handle these stencils in a global manner, we follow the concept of *impact zone* (IZ) of [Bridson et al. 2002; Harmon et al. 2008]. As each constraint gradient depends on four vertices, an impact zone is defined as a set of constraint gradients with overlapping stencils. The projection operation (Equ. 3) is performed per IZ, and each IZ corresponds to one linear system of Equ. 7.

5 Diffusion of Displacement Vectors

Cloth is often regarded as continuum material, and the dynamics model is designed to demonstrate corresponding behavior. In computer simulation particles are inter-connected by springs, and their motion is subject to physics laws. Thanks to the implicit dynamics solver being used, the particle velocities form a smoothly varying vector field. That is to say if two particles are topologically very close, so are their respective velocities. However when collision constraints enforce position or velocity adjustment, discontinuities will be introduced to the smooth vector field, causing not only visual artifacts but also simulation instability. Our solution is to propagate the correcting displacement vectors from the penetrating vertices to the vicinity, creating a smooth vector field of corrections to be applied globally. We use the constrained diffusion method given in [Xu et al. 2009b; Xu et al. 2009a] to compute harmonic vector field.

6 Applications and Results

To test the effectiveness of proposed method, we designed several experiments. In each case at least one mesh is specified as oriented

surface, so the mesh-mesh collision is handled with the proposed method. The self-collision of each deformable mesh is handled with the traditional CCD-based method.

The first experiment (Fig. 1a) simulates a spike (represented as an oriented surface) stabbing a sheet (18,896 vertices), to show our method's ability to process sharp geometry. At the beginning, the collision response is disabled and the sheet is pierced by the spike on purpose. A while later our DCD-based collision response is enabled, and the existing penetration gets resolved and no new penetration occurs thereafter.

The second experiment (Fig. 1b) simulates two tori colliding with each other. Both tori are modeled as oriented surface meshes of 576 vertices. Air pressure is computed based on its volume and applied to vertices. We vary the post-response distance $\mathbf{d}$ in Equ. 8 to emulate behaviors from elastic to inelastic collisions. Note the red torus is not subject to gravity (consider it as aired balloon).

In the third experiment (Fig. 4), the proposed method has been integrated into a cloth simulator based on [Narain et al. 2012]. While simulating a dressed dummy, body parts are modeled as oriented surface meshes. The self-collision of the garment mesh is handled with the simulator's built-in CCD pipeline, while the body/garment collision is processed with our DCD-based algorithm. The simulator runs at time step $\Delta t = 0.005s$, and outputs one frame data every eight steps. To demonstrate that our method tolerates much larger time step than CCD, the body/garment collision is handled once for every eight steps – only at the steps when frame data is to be output! The accompanying video shows there is no quality loss in our simulation result. Moreover, the new simulator is faster than the old one – it cuts the overall collision handling time by 40%.

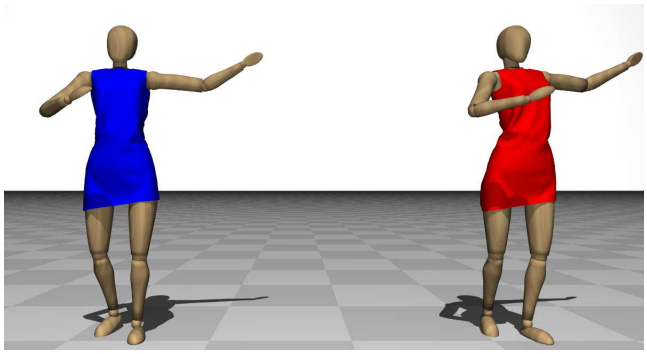

Figure 4: *By taking a much larger time step to do solid/cloth collision faster, our result (right) does not suffer any quality loss compared to Narain's (left).*

7 Discussions and Conclusion

We have presented a new collision response approach to untangle existing penetrations. In many circumstances where surface orientation is either explicitly given or implied in the context (e.g. closed surfaces), this method is a competitive alternative to the popular CCD-based approach.

Self-healing is the primary advantage of DCD-based response method. Even if in certain steps there are detected collisions that can not be resolved right away due to various reasons, they will not ruin the simulation in the following steps. Of course, the penetration will be processed at earliest possible time.

Another significant benefit of DCD-based approach is the large time step. As physically based simulation is often computational expensive, acceleration is pursued in two aspects: speeding-up the per-step computation and using larger time steps. As new technique [Liu et al. 2013] has made it possible for the dynamics solver to take large time step, the CCD-based collision handling, which is sensitive to the time step, becomes the bottleneck of the simulation. Our DCD-based approach provides a solution to this issue. In real production, the artist can opt, in some designated steps, for incomplete response, as long as the intersecting regions do not expand too much and the rendered images do not suffer from visual artifacts. In video games, insignificant penetrations are often acceptable as long as they do not create obvious artifacts.

Limitations. For the DCD-based response strategy to be widely adopted, it must be extended to be able to handle un-oriented surfaces. As the algorithm relies on the front/back identification, we must, at least, be able to dynamically identify the faces of the colliding regions. This can be achieved either by global intersection analysis [Baraff et al. 2003] or from history information. In the latter case, the history will only be used to identify front/back, but no cubic collision solver should be involved.

Although diffusion field is created to propagate the positional displacement, we noticed sometimes the relocated positions introduce self-collisions in the deformable mesh. It is desirable that resolving one collision does not generate a new one.

Acknowledgements

This work was supported by xxxx under the grants #yyyy.

References

AINSLEY, S., VOUGA, E., GRINSPUN, E., , AND TAMSTORF, R. 2012. Speculative parallel asynchronous contact mechanics. *ACM Trans. Graph. 31*, 6, 151:1–151:8.

BARAFF, D., WITKIN, A., AND KASS, M. 2003. Untangling cloth. *ACM Trans. Graph. 22*, 3, 862–870.

BRIDSON, R., FEDKIW, R., AND ANDERSON, J. 2002. Robust treatment of collisions, contact and friction for cloth animation. *ACM Trans. Graph. 21*, 3, 594–603.

BROCHU, T., EDWARDS, E., AND BRIDSON, R. 2012. Efficient geometrically exact continuous collision detection. *ACM Trans. Graph. 31*, 4, 96:1–96:7.

HARMON, D., VOUGA, E., TAMSTORF, R., AND GRINSPUN, E. 2008. Robust treatment of simultaneous collisions. *ACM Trans. Graph. 27*, 3.

HARMON, D., VOUGA, E., SMITH, B., TAMSTORF, R., AND GRINSPUN, E. 2009. Asynchronous Contact Mechanics. *ACM Trans. Graph. 28*, 3.

LIU, T., BARGTEIL, A. W., O'BRIEN, J. F., AND KAVAN, L. 2013. Fast simulation of mass-spring systems. *ACM Trans. Graph. 32*, 6, 1 – 7.

NARAIN, R., SAMII, A., AND O'BRIEN, J. F. 2012. Adaptive anisotropic remeshing for cloth simulation. *ACM Transactions on Graphics 31*, 6, 147:1–10. Proceedings of ACM SIGGRAPH Asia 2012, Singapore.

PROVOT, X. 1997. Collision and self-collision handling in cloth model dedicated to design garments. In *EG Workshop on Computer Animation and Simulation*, 177–189.

VOLINO, P., AND MAGNENAT-THALMANN, N. 2006. Resolving surface collisions through intersection contour minimization. *ACM Trans. Graph. 25*, 3, 1154–1159.

WICKE, M., LANKER, H., AND GROSS, M. 2006. Untangling cloth with boundaries. In *Proc. of Vision, Modeling, and Visualization*, 349–356.

XU, K., COHEN-OR, D., JU, T., LIU, L., ZHANG, H., ZHOU, S., AND XIONG, Y. 2009. Feature-aligned shape texturing. *ACM Trans. Graph. 28*, 5, 108:1–108:7.

XU, K., ZHANG, H., COHEN-OR, D., AND XIONG, Y. 2009. Dynamic harmonic fields for surface processing. *Computers & Graphics 33*, 3, 391–398.

YE, J., AND ZHAO, J. 2012. The intersection contour minimization method for untangling oriented deformable surfaces. In *Proc. Symp. Computer Animation*, 311–316.